arXiv:1607.07472v1 [cs.RO] 25 Jul 2016

Efficient Multi-Agent Global Navigation Using Interpolating Bridges

Liang He[1] and Jia Pan[2] and Dinesh Manocha[1]

[1] UNC Chapel Hill
[2] City University of Hong Kong

Abstract. We present a novel approach for collision-free global navigation for continuous-time multi-agent systems with general linear dynamics. Our approach is general and can be used to perform collision-free navigation in 2D and 3D workspaces with narrow passages and crowded regions. As part of pre-computation, we compute multiple bridges in the narrow or tight regions in the workspace using kinodynamic RRT algorithms. Our bridge has certain geometric characteristics, that enable us to calculate a collision-free trajectory for each agent using simple interpolation at runtime. Moreover, we combine interpolated bridge trajectories with local multi-agent navigation algorithms to compute global collision-free paths for each agent. The overall approach combines the performance benefits of coupled multi-agent algorithms with the precomputed trajectories of the bridges to handle challenging scenarios. In practice, our approach can handle tens to hundreds of agents in real-time on a single CPU core in 2D and 3D workspaces.

1 Introduction

Multi-agent navigation algorithms are widely used for motion planning among static and dynamic obstacles. The underlying applications include cooperative surveillance, sensor networks, swarm navigation, and simulation of animated characters or human crowds in games and virtual worlds. One key problem in multi-agent navigation is the computation of collision-free trajectories for agents, given their own initial and goal positions.

This problem has been studied extensively in robotics, AI, and computer animation. At a broad level, prior approaches can be classified into *coupled* or *decoupled* planners. A coupled planner aggregates all the individual robots into one large composite system, and leverages classical motion planners (e.g., sampling-based planners) to compute collision-free trajectories for all agents. On the other hand, decoupled planners computes a trajectory for each robot individually for a short horizon (e.g., a few time-steps), and then performs a velocity coordination to resolve the collision between the local trajectories of all agents. Different techniques have been proposed to compute local collision-free paths or schedule their motion.

Coupled planners are (probabilistically) complete in theory and thus can provide rigorous guarantees about collision avoidance between the agents and

the obstacles. However, as the number of agents in the scene increases, the resulting dimension of the system's configuration space increases linearly. As a result, current coupled algorithms are only practical for a few agents. Decoupled planners are generally faster because fewer degrees of freedom are taken into account at a time. Unfortunately, they are usually not complete because, the velocity coordination may not resolve all collisions, and the agents may get stuck in crowded scenarios or block each other (e.g., due to the inevitable collision states [1,2]). Thus, in challenging scenarios with narrow passages, a decoupled planner may take a long time or even be unable to find a solution when one exists. In addition, the velocity coordination can be slow in crowded environments, where all agents have to move toward their goals at very small steps.

Main Results: In this paper, we present a novel method to accelerate the performance of decoupled multi-agent planners in crowded or challenging environments. Our approach is limited to scenarios with static obstacles or dynamic obstacles whose trajectories are kn own beforehand. The key idea is to compute bridges in the narrow passages or challenging areas of the workspace, which are collision-free regions of the workspace and have certain geometric navigation characteristics. As an agent approaches the bridge, we use the precomputed local trajectories associated with that bridge to guide the agent toward the goal position. We also present an efficient scheduling scheme which enables multiple agents to share a single bridge efficiently. The overall trajectory of each agent is calculated by using an optimal trajectory generation algorithm [3] along the interpolated path in the bridge, which combines the efficiency of the decoupled methods with the completeness of the interpolating bridges.

A novel component of our approach is the computation of interpolating bridges in 2D and 3D workspaces. Each bridge lies in the collision-free space, and its boundaries are calculated using kinodynamic RRT algorithms. Our approach guarantees that when an agent enters a bridge with a velocity satisfying suitable conditions (as later will be discussed in detail), it can always compute a collision-free trajectory that lies within the bridge. Furthermore, we present an inter-trajectory scheduling scheme for multiple agents sharing a bridge for navigation that has a small runtime overhead. We present efficient algorithms to compute these bridges in 2D and 3D workspaces, to generate trajectories within bridges at runtime, and to schedule agents sharing the bridges for efficient global navigation. We highlight the performance of our method in several challenging 2D and 3D workspsaces with narrow passages. We demonstrate its performance by comparing it with prior multi-agent navigation approaches based on local collision avoidance. In practice, our method can handle about 50 agents in real-time on a single CPU core for both 2D and 3D scenarios.

The rest of the paper is organized as follows. We give a brief survey of prior work in Section 2. We introduce our notation and provide an overview of our approach in Section 3. We describe the bridge computation and global multi-agent navigation algorithm in Section 4 and Section 6, respectively. We analyze our method's properties in Section 7, and finally highlight its performance in Section 8.

2 Related Work

There has been extensive work on multi-agent motion planning. This work includes many reactive methods such as RVO [5], HRVO [6], and their variants. These techniques compute a feasible movement for each agent such that it can avoid other agents and obstacles in a short time horizon. However, they cannot provide mathematical guarantees about whether or not agents can always find collision-free trajectories. In particular, they may not be able to avoid the inevitable collision states (ICS) [1,2,7,8] in the configuration space, due to robots' dynamical constraints or obstacles in the scenario. Some methods [8,9,10] provide partial solutions to these problems, but they still cannot guarantee avoidance of all ICS in the long horizon while working in a crowd scenario with narrow passages. Actually, even for a scenario without static obstacles, it is still difficult to achieve robust collision-avoidance coordination when there is a large number of agents [11,12].

The simplest solution to the difficulty of inevitable collision states is to design suitable protocols for multi-robot coordination/interaction [13,14]. Some other approaches precompute roadmaps or corridors in the entire workspace to achieve high-quality path planning [15,16]. However, these methods are not complete and may provide sub-optimal trajectories. Our method also leverages pre-computed bridges to deal with the navigation challenges in narrow or crowded regions. However, the bridges used in our approach have special properties beneficial for efficient global navigation in challenging areas in the workspace.

Centralized multi-agent navigation approaches usually leverage global single-robot planning algorithms (such as PRM or RRT) to compute a roadmap or grids for the high-level coordination [17,18]. Compared to the decentralized methods, these algorithms compute all agents' trajectories simultaneously and thus can better handle the complex interactions among agents. These methods can also be extended to handle non-holonomic multi-agent systems (e.g., systems composed of differential-drive robots), by using local planners like RRT [19], RRT* [20], or other algorithms that can deal with differential dynamics [3]. In addition to their benefit in terms of finding feasible trajectories, the centralized algorithms can avoid deadlock cases by leveraging high-level scheduling or coordination strategies, either coupled [21,22] or decoupled [23,24,25,26]. In general, these strategies only work in theory, because they have to ignore the robot's dynamics and assume the robots to be operating in a discrete state space. Finally, centralized multi-agent navigation algorithms are computationally expensive and can sometimes be too slow for real-world applications.

3 Overview

3.1 Problem Definition

Our goal is to enable a group of agents to reach their individual goals in a safe and efficient manner. Our approach is designed for continuous-time multi-agent systems with general linear dynamics. During the navigation, the agents should

avoid collisions with static obstacles in the environment as well as with other agents.

This problem can be formally defined as follows. We take as given a set of n decision-making agents sharing an environment consisting of obstacles. For simplicity, we assume the geometric shape of each agent is represented as a disc of radius r, and its current position and velocity are denoted as $\mathbf{p}$ and $\mathbf{v}$, respectively. We also assume kinodynamic constraints on the agent's motion: $\|\mathbf{v}\| \leq v_{\max}$ and $\|\mathbf{a}\| \leq a_{\max}$, where $v_{\max}$ and $a_{\max}$ are the maximum allowed velocity and acceleration for the agent's velocity $\mathbf{v}$ and acceleration $\mathbf{a}$. The task of each agent is to compute a trajectory toward its goal position G from its initial position I. The trajectory should be collision-free and also satisfy other dynamics constraints.

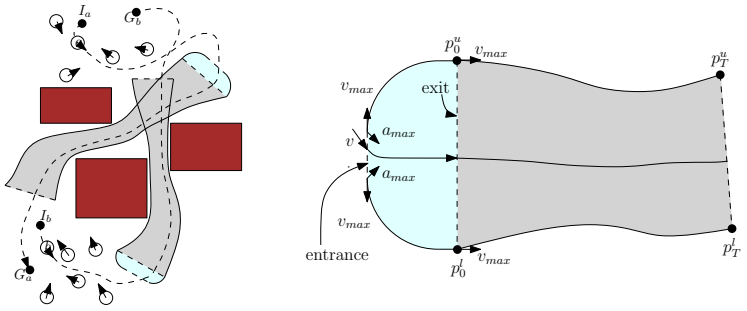

Fig. 1: An overview of our approach: The left figure shows a scene with three obstacles (the brown boxes), agents (the circles) want to reach their individual goals $G_a, G_b, ...$ from their initial positions $I_a, I_b, ...$ We use two bridges (the gray regions) in the environment to help the agents navigate through the crowded areas efficiently. The agent a first goes toward the entrance (the cyan region) of one bridge following a path from I_a to the entrance. Next, a goes inside the entrance where its velocity is gradually adjusted before entering the bridge. Finally, a goes through the bridge and arrives at its goal G_a. The agent b arrives at its goal G_b through another bridge. The dashed line denotes an agent's trajectory. These bridges enable global navigation with safety guarantees. The right figure shows the entrance to a bridge: Given an agent with velocity $\mathbf{v}$, we use the entrance to gradually change the agent's velocity so that while entering the bridge the agent's velocity is equal to the initial velocity of the bridge's boundary trajectories.

3.2 Our Approach

Given a crowded scenario as shown in Figure 1, the most efficient navigation strategy for a group of agents should be as follows: in the open space outside the narrow passage, each agent can navigate according to a path computed by

optimal motion planning approaches; in the narrow area, agents should maintain a line or some formation, and then pass through the narrow region in some order.

For this purpose, we allocate a set of guidance channels called *bridges* in the narrow parts of the workspace, as shown in Figure 1. Each bridge lies totally in the collision-free subset of the configuration space and its boundary is computed using RRT-based collision-free paths. Furthermore, each bridge has an attractive characteristic such that for an agent entering the bridge with its velocity satisfying a special constraint, we can always compute a trajectory that completely lies inside the bridge. In order to pass through the narrow passages of the workspace, an agent will first move toward one of the bridges according to its own dynamics. However, the agent may arrive at the bridge with an arbitrary velocity. Thus, the agent must first enter an area called *entrance* in front of the bridge. The entrance will gradually adjust the agent's velocity to make sure it satisfies the requirements with respect to that bridge. Once the agent leaves the entrance and enters inside the bridge, it can follow a collision-free trajectory to pass through the crowded or narrow region efficiently and safely. After leaving the bridge, the agent can switch back to the local coordination strategy and move toward its goal. In this way, the bridges can be viewed as a "highway system" for the agents to efficiently travel through challenging scenarios.

While we can build one bridge for each agent, this may result in many overlaps between the bridges, and therefore the agents will need to slow down or even wait while moving along the bridges. A better solution is to build a few bridges that can be reached by all agents, and allow multiple agents to share one bridge by using suitable scheduling schemes.

4 Bridge Construction

In this section, we describe the details of how to automatically compute the bridges in the workspace. We will first describe our approach for a 2D workspace, which will later be extended to 3D workspaces.

We start from a zero-width bridge, which is a collision-free trajectory connecting a pair of start and goal positions in the crowded region. The trajectory is computed using a kinodynamic RRT planner [19], and has an initial velocity of $\mathbf{v}_0$ with a magnitude $v_{\max}$. We then incrementally enlarge the bridge's width until it touches one obstacle in the scenario. The built bridge has a beneficial property alllowing that as long as an agent enters the bridge with the velocity $\mathbf{v}_0$, it can always use an efficient interpolation scheme to calculate a collision-free trajectory passing through the bridge. However, an agent may enter a bridge with an arbitrary velocity and thus violate the bridge's constraint on the entering velocity. To solve this problem, we add an entrance region in front of the bridge, which provides sufficient room for an agent to gradually adjust its velocity toward $\mathbf{v}_0$ in order to leverage the bridge for safe and efficient navigation. Examples of 2D bridges and their entrances are shown in Figures 1, 2. Finally, for each entrance we compute a backward-reachable set, i.e., the set of positions from which an agent can reach the entrance. We also compute a

forward-reachable set for the end line of the bridge, i.e., the set of positions that can be reached by an agent coming out of the bridge. All agents with their start positions inside the backward-reachable set and goal positions inside the forward-reachable set will leverage the constructed bridge for navigation. We repeatedly add more bridges for the remaining agents until each agent has one associated bridge.

4.1 Iterative Bridge Enlargement

We first use the kinodynamic RRT planner [19] to connect a pair of starting and goal positions in the crowded scenarios, and the result is a trajectory $\{\mathbf{p}_0, ..., \mathbf{p}_T\}$. We consider the trajectory as a bridge with zero width, and its two boundary trajectories are $\{\mathbf{p}_0^u, ..., \mathbf{p}_T^u\}$ and $\{\mathbf{p}_0^l, ..., \mathbf{p}_T^l\}$ respectively, where $\mathbf{p}_i = \mathbf{p}_i^u = \mathbf{p}_i^l$ are overlapped points. Starting from this initial bridge, we incrementally enlarge the bridge's width until it hits one obstacle. The algorithm is as shown in Algorithm 1.

Algorithm 1: Construct a 2D bridge

input : Start and goal points $\mathbf{p}_0$ and $\mathbf{p}_T$
output: A valid bridge

`/* Initialize with a zero-width bridge */`
$\{\mathbf{p}_0, ..., \mathbf{p}_T\} \leftarrow RRT(\mathbf{p}_0, \mathbf{p}_T)$
`/* Set the overlapped bridge upper and lower boundaries */`
$\{\mathbf{p}_0^u, ..., \mathbf{p}_T^u\} \leftarrow \{\mathbf{p}_0, ..., \mathbf{p}_T\}$
$\{\mathbf{p}_0^l, ..., \mathbf{p}_T^l\} \leftarrow \{\mathbf{p}_0, ..., \mathbf{p}_T\}$
while *true* **do**
  `/* Check collision for the bridge area between two boundaries */`
  collision $\leftarrow$ bridgeCollide($\{\mathbf{p}_0^u, ..., \mathbf{p}_T^u\}, \{\mathbf{p}_0^l, ..., \mathbf{p}_T^l\}$)
  if *collision* **then**
    break
  `/* Change the start and goal positions of two boundaries */`
  $\mathbf{p}_0^u \leftarrow \mathbf{p}_0^u + \Delta\mathbf{p}_0$, $\mathbf{p}_T^u \leftarrow \mathbf{p}_T^u + \Delta\mathbf{p}_T$
  $\mathbf{p}_0^l \leftarrow \mathbf{p}_0^l - \Delta\mathbf{p}_0$, $\mathbf{p}_T^l \leftarrow \mathbf{p}_T^l - \Delta\mathbf{p}_T$
  `/* Generate new bridge boundaries using RRT */`
  $\{\mathbf{p}_0^u, ..., \mathbf{p}_T^u\} \leftarrow$RRT($\mathbf{p}_0^u, \mathbf{p}_T^u$)
  $\{\mathbf{p}_0^l, ..., \mathbf{p}_T^l\} \leftarrow$RRT($\mathbf{p}_0^l, \mathbf{p}_T^l$)
return A bridge b with boundaries $\{\mathbf{p}_0^u, ..., \mathbf{p}_T^u\}$ and $\{\mathbf{p}_0^l, ..., \mathbf{p}_T^l\}$

4.2 Trajectory Generation in a Bridge

Given a 2D bridge, as shown in Figure 2, we can show that if the agent enters the bridge with a velocity equal to the initial velocities of the bridge's two boundary trajectories, we can use an efficient interpolation scheme to generate a trajectory

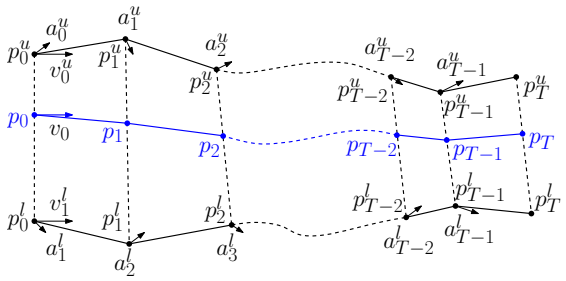

Fig. 2: Trajectory interpolation in a 2D bridge: The bridge is bounded by an upper trajectory $\{\mathbf{p}_0^u, \mathbf{p}_1^u, ..., \mathbf{p}_{T-1}^u, \mathbf{p}_T^u\}$ and a lower trajectory $\{\mathbf{p}_0^l, \mathbf{p}_1^l, ..., \mathbf{p}_{T-1}^l, \mathbf{p}_T^l\}$, where the velocity and acceleration at each trajectory point $\mathbf{p}_i$ are $\mathbf{v}_i$ and $\mathbf{a}_i$ respectively. The start line and end line of the bridge are $\mathbf{p}_0^u\mathbf{p}_0^l$ and $\mathbf{p}_T^u\mathbf{p}_T^l$ respectively. The agent enters the bridge at the position $\mathbf{p}_0$ on the start line with a velocity $\mathbf{v}_0$. If $\mathbf{v}_0 = \mathbf{v}_0^u = \mathbf{v}_0^l$, the trajectory interpolation scheme can compute the agent's trajectory as $\{\mathbf{p}_0, ..., \mathbf{p}_T\}$ which will be located completely inside the bridge.

for the agent such that the agent will always stay inside the bridge and compute a path that is collision-free with all static obstacles.

The interpolation details are described in Algorithm 2. In particular, we choose the acceleration of the agent at each time step $i\Delta t$ as a linear interpolation of the accelerations of the corresponding waypoints on the boundary trajectories: $\mathbf{a}_i = (1-r)\mathbf{a}_i^u + r\mathbf{a}_i^l$. The interpolation coefficient r is the ratio based on which the agent's initial position $\mathbf{p}_0$ partitions the bridge area's start line $\mathbf{p}_0^u\mathbf{p}_0^l$: $r = \frac{|\mathbf{p}_0^u\mathbf{p}_0|}{|\mathbf{p}_0\mathbf{p}_0^l|}$. We can prove that in this way, the generated trajectory $\{\mathbf{p}_0, ..., \mathbf{p}_T\}$ always stays inside the bridge:

Theorem 1. The interpolated trajectory $\{\mathbf{p}_0, ..., \mathbf{p}_T\}$ always stays inside the bridge.

Proof. We first show that the interpolated trajectory has the following two properties $\mathbf{p}_i = (1-r)\mathbf{p}_i^u + r\mathbf{p}_i^l$ and $\mathbf{v}_i = (1-r)\mathbf{v}_i^u + r\mathbf{v}_i^l$, for $i = 1, ..., T$. We can prove these two statements by induction on i. If $i = 1$, $\mathbf{p}_0 = (1-r)\mathbf{p}_0^u + r\mathbf{p}_0^l$ is trivial because this is the definition of r; and since $\mathbf{v}_0 = \mathbf{v}_0^u = \mathbf{v}_0^l$ as required by the bridge's definition, $\mathbf{v}_0 = (1-r)\mathbf{v}_0^u + r\mathbf{v}_0^l$ is also obvious. Now consider the case in which $i > 1$:

$$\begin{aligned}\mathbf{v}_i &= \mathbf{v}_{i-1} + \mathbf{a}_{i-1}\Delta t = (1-r)\mathbf{v}_{i-1}^u + r\mathbf{v}_{i-1}^l + [(1-r)\mathbf{a}_i^u + r\mathbf{a}_i^l]\Delta t \\ &= (1-r)[\mathbf{v}_{i-1}^u + \mathbf{a}_i^u\Delta t] + r[\mathbf{v}_{i-1}^l + \mathbf{a}_i^l\Delta t] = (1-r)\mathbf{v}_i^u + r\mathbf{v}_i^l.\end{aligned} \tag{1}$$

Similarly, we have

$$\mathbf{p}_i = \mathbf{p}_{i-1} + \mathbf{v}_{i-1}\Delta t + \frac{1}{2}\mathbf{a}_{i-1}(\Delta t)^2 = (1-r)\mathbf{p}_i^u + r\mathbf{p}_i^l. \tag{2}$$

Based on the induction hypothesis, the waypoints of the interpolated trajectory will always be a linear interpolation of the corresponding waypoints of the

bridge's two boundary trajectories. Thus, the interpolated trajectory must be contained inside the bridge.

Algorithm 2: Trajectory generation in a 2D bridge

input : The bridge's two boundary trajectories $\{\mathbf{p}_0^u, ..., \mathbf{p}_T^u\}$ and $\{\mathbf{p}_0^l, ..., \mathbf{p}_T^l\}$, along with each waypoint $\mathbf{p}_i$'s velocity $\mathbf{v}_i$ and acceleration $\mathbf{a}_i$. The initial position $\mathbf{p}_O$ and velocity $\mathbf{v}_0$ when the agent enters the bridge.

output: The agent's trajectory $\{\mathbf{p}_0, ..., \mathbf{p}_T\}$

`/* Compute the ratio into which` $\mathbf{p}_0$ `partitions the start line` $\mathbf{p}_0^u\mathbf{p}_0^l$ `*/`

$r = \frac{|\mathbf{p}_0^u\mathbf{p}_0|}{|\mathbf{p}_0\mathbf{p}_0^l|}$;

for $i = 1, ..., T-1$ **do**

  $\mathbf{a}_i = (1-r)\mathbf{a}_i^u + r\mathbf{a}_i^l$;

  $\mathbf{p}_{i+1} = \mathbf{p}_i + \mathbf{v}_i \Delta t + \frac{1}{2}\mathbf{a}_i(\Delta t)^2$;

  $\mathbf{v}_{i+1} = \mathbf{v}_i + \mathbf{a}_i \Delta t$;

4.3 2D Entrance Construction and Trajectory Generation

When an agent enters the bridge, its velocity must be the same as the initial velocity of the bridge's boundary trajectories. However, agents may arrive at the bridge with an arbitrary velocity, and thus the bridge may not be able to generate a trajectory for an agent that is completely inside the bridge. Our solution is to leave some space in front of the bridge called the *entrance* to the bridge. This space works like a buffer zone where the agent can gradually adjust its velocity to meet the bridge's requirement.

An example of the 2D entrance is shown in the right side in Figure 1, which includes four parts: one start line, one end line, and two boundary curves. The end line of the entrance is also the start line of the bridge. Each boundary curve is composed of a parabola followed by a line. The two parabolas correspond to the trajectories in which the agent's velocity is gradually rotated by 90 degrees from the initial $\pm v_{\max}\hat{\mathbf{y}}$ to the final $v_{\max}\hat{\mathbf{x}}$. The velocity change is achieved by using a constant acceleration $\mathbf{a}$, the magnitude of which is $a_{\max}$ and the direction of which is $\frac{\sqrt{2}}{2}\hat{\mathbf{x}} \mp \frac{\sqrt{2}}{2}\hat{\mathbf{x}}$. For convenience, we assume the bridge's start line is along the $\hat{\mathbf{y}}$ direction. The two lines following the parabolas have a length $(\sqrt{\frac{2}{3}} - \sqrt{\frac{1}{2}})\frac{v_{\max}^2}{a_{\max}}$ each.

The shape of the entrance chosen has the property that, for any agent with an arbitrary velocity $\mathbf{v}$ ($|\mathbf{v}| \leq v_{\max}$) entering the entrance, we can always find a sequence of accelerations to gradually change its velocity from $\mathbf{v}$ to $v_{\max}\hat{\mathbf{x}}$(i.e., with a magnitude of $v_{\max}$ and perpendicular to the bridge's start line). After this velocity adjustment, the agent will have a velocity equal to the initial velocity of the bridge's boundary trajectories when entering the bridge. According to

Section 4.2, this property is required for the trajectory interpolation inside the bridge.

We can choose the acceleration such that, during the adjustment of the velocity, the agent will always stay inside the area of the entrance and is therefore, guaranteed to avoid collisions with other agents and obstacles. In algorithm 2, we provide the details about velocity adjustment. In particular, given an agent entering the entrance at a velocity $\mathbf{v} = (v_x, v_y)$, our goal is to increase v_x to $v_{\max}$ and decrease v_y to 0. We first determine a suitable acceleration (a_x, a_y) by comparing the speed gaps in both directions (line 1). This acceleration guarantees that v_x will increase to $v_{\max}$ after v_y arrives at 0, which is necessary to keep the agent inside the entrance region. After v_y becomes 0, the acceleration along the $\hat{\mathbf{y}}$ will become zero, while the acceleration along $\hat{\mathbf{x}}$ remains the same. The agent moves on until its velocity is $v_{\max}\hat{\mathbf{x}}$. After that, if the agent has not reached the exit line, it will continue with the current velocity. After computing the extreme positions at which the resulting trajectory will arrive, we can show that the trajectory will never go out of the entrance region and thus is guaranteed to be collision-free with static obstacles.

4.4 Bridges' Forward and Backward Reachable Regions

After generating the bridge and its entrance, we then compute its forward and backward reachable regions. We denote $\mathcal{R}^+[\mathbf{p}, \tau]$ as the forward-reachable set of a position $\mathbf{p}$, (i.e., the set of positions that can be reached from $\mathbf{p}$ with time less than τ), and denote $\mathcal{R}^-[\mathbf{p}, \tau]$ as the backward-reachable set, (i.e., the set of positions that can reach $\mathbf{p}$ with time less than τ):

$$\mathcal{R}^+(\mathbf{p}, \tau) = \{\mathbf{p}' | \mathit{time}(\mathbf{p}, \mathbf{p}') < \tau\} \tag{3}$$

$$\mathcal{R}^-(\mathbf{p}, \tau) = \{\mathbf{p}' | \mathit{time}(\mathbf{p}', \mathbf{p}) < \tau\}, \tag{4}$$

where $\mathit{time}(\mathbf{x}, \mathbf{y})$ measures the time an agent needs to move from a starting position $\mathbf{x}$ to a goal position $\mathbf{y}$. Both reachability sets can be efficiently estimated using the method proposed in [3].

Leveraging the concept of forward and backward reachable sets, we can compute the forward and backward reachable regions for the entire bridge as

$$\mathcal{R}^+(\mathit{bridge}, \tau) = \bigcup_{\mathbf{p} \in \mathit{bridge\ end\ line}} \mathcal{R}^+(\mathbf{p}, \tau) \tag{5}$$

$$\mathcal{R}^-(\mathit{bridge}, \tau) = \bigcup_{\mathbf{p} \in \mathit{entrance\ start\ line}} \mathcal{R}^-(\mathbf{p}, \tau). \tag{6}$$

4.5 Bridge Assignment among Agents

We have finished describing how to build a bridge, its entrance, and the corresponding reachability regions. We will now discuss how to build a set of bridges for all the agents in the scenario, as well as how to assign a bridge to each agent.

Our solution is to first build a bridge for a single agent. Next, we check whether or not there are any other agents whose initial and goal positions are within the backward and forward reachable regions of this bridge. If so, these agents can also leverage this bridge for navigation, and thus we assign this bridge to these agents. We repeatedly add more bridges until all agents have one associated bridge. The algorithm is as shown in Algorithm 4.

Algorithm 3: Generating bridges for all agents

input : Initial and goal configurations pairs $IG = \{(I_i, G_i)\}_{i=0}^{n-1}$ for all n agents
output: A set of bridges $B = \{B_j\}$

$B \leftarrow \emptyset$
`/* compute the bridges for all of the agents */`
while $IG \neq \emptyset$ **do**
  Select two configurations (I, G) from the set IG
  Construct a 2D bridge b and its entrance using Algorithm 1
  Compute the reachability regions $\mathcal{R}^+(b, \tau)$ and $\mathcal{R}^-(b, \tau)$
  `/* compute the bridge for the agents */`
  for $(I_i, G_i) \in IG$ **do**
    if $I_i \in \mathcal{R}^-(b, \tau)$ and $G_i \in \mathcal{R}^+(b, \tau)$ **then**
      Remove (I_i, G_i) from IG
      Assign bridge b to the agent i
      Add bridge b to B
return B

4.6 3D Bridge and Entrance

The bridge and entrance algorithms described above can be extended to the 3D workspace. Figure 3 illustrates the 3D bridge and the corresponding trajectory interpolation algorithm. The boundary of the 3D bridge is composed of K trajectories $\{\mathbf{p}_0^k, \mathbf{p}_1^k, ..., \mathbf{p}_{T-1}^k, \mathbf{p}_T^k\}$, computed by the RRT algorithm, where $k = 1, ..., K$. The initial points $\mathbf{p}_0^k$ of all these K trajectories are located on the same plane(called the start plane of the bridge), and their initial velocities $\mathbf{v}_0^k$ are the same, all having the magnitude of $v_{\max}$. Given an agent entering the bridge at position $\mathbf{p}_0$ on the start plane, we can compute its trajectory using interpolation as follows. First, from all the initial points, we can select three points whose convex combination can be used for the point $\mathbf{p}_0$. W.l.o.g., we assume that these three points are $\mathbf{p}_0^0, \mathbf{p}_0^1, \mathbf{p}_0^2$, and their convex combination is $\mathbf{p}_0 = u\mathbf{p}_0^0 + v\mathbf{p}_0^1 + w\mathbf{p}_0^2$, where $0 \leq u, v, w \leq 1$ and $u + v + w = 1$. Similar to Theorem 1, we can show that if we choose the acceleration at time t to be $\mathbf{a}_t = u\mathbf{a}_t^0 + v\mathbf{a}_t^1 + w\mathbf{a}_t^2$, the resulting trajectory will have velocity $\mathbf{v}_t = u\mathbf{v}_t^0 + v\mathbf{v}_t^1 + w\mathbf{v}_t^2$ and position $\mathbf{p}_t = u\mathbf{p}_t^0 + v\mathbf{p}_t^1 + w\mathbf{p}_t^2$ at time t. It follows that this trajectory will always stay inside the 3D bridge. The 3D entrance construction and trajectory generation

algorithms are also similar to the 2D case, except that the entrance's boundary is now a surface bounded by a series of K parabolas.

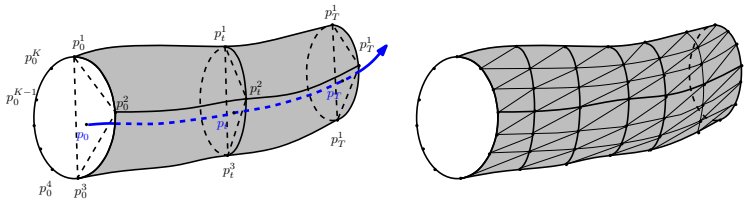

Fig. 3: The left figure shows the trajectory interpolation in a 3D bridge: the bridge is bounded by a set of boundary trajectories $\{\mathbf{p}_0^k, \mathbf{p}_1^k, ..., \mathbf{p}_{T-1}^k, \mathbf{p}_T^k\}$, where $k = 1, ..., K$. The agent enters the bridge at the position $\mathbf{p}_0$ on the start plane. The trajectory interpolation is performed by first choosing three points from $\{\mathbf{p}_i^k\}$ that can provide a convex combination for the point $\mathbf{p}_0$. Here we assume the first three points are chosen, and thus $\mathbf{p}_0 = u\mathbf{p}_0^1 + v\mathbf{p}_0^2 + w\mathbf{p}_0^3$, where $0 \leq u, v, w \leq 1$ and $u + v + w = 1$. Similar to the 2D bridge, we can always choose a suitable acceleration $\mathbf{a}_t$ so that the resulting trajectory $\{\mathbf{p}_0, ..., \mathbf{p}_T\}$ completely locates inside the bridge. The right figure shows the triangulation for the 3D bridge shown in the left figure. In this case, each point is one waypoint in one of the K boundary trajectories.

4.7 Collision Checking for a Bridge

Given a bridge, we need to make sure that it is collision-free. To perform collision checking on a bridge, we triangulate the bridge boundary, and then perform collision checking between the triangulated bridge boundary with static obstacles in the environment. The left side in figure 3 shows the triangulation result for the 3D bridge . The collision checking is performed using traditional bounding volume techniques [27].

5 Bridges' Forward and Backward Reachable Regions

After generating the bridge and its entrance, we then compute its forward and backward reachable regions. We denote $\mathcal{R}^+[\mathbf{p}, \tau]$ as the forward-reachable set of a position $\mathbf{p}$,(i.e., the set of positions that can be reached from $\mathbf{p}$ with time less than τ), and denote $\mathcal{R}^-[\mathbf{p}, \tau]$ as the backward-reachable set, (i.e., the set of positions that can reach $\mathbf{p}$ with time less than τ):

$$\mathcal{R}^+(\mathbf{p}, \tau) = \{\mathbf{p}' | time(\mathbf{p}, \mathbf{p}') < \tau\} \tag{7}$$

$$\mathcal{R}^-(\mathbf{p}, \tau) = \{\mathbf{p}' | time(\mathbf{p}', \mathbf{p}) < \tau\}, \tag{8}$$

where $time(\mathbf{x}, \mathbf{y})$ measures the time an agent needs to move from a starting position $\mathbf{x}$ to a goal position $\mathbf{y}$. Both reachability sets can be efficiently estimated using the method proposed in [3].

Leveraging the concept of forward and backward reachable sets, we can compute the forward and backward reachable regions for the entire bridge as

$$\mathcal{R}^+(bridge, \tau) = \bigcup_{\mathbf{p} \in bridge\ end\ line} \mathcal{R}^+(\mathbf{p}, \tau) \tag{9}$$

$$\mathcal{R}^-(bridge, \tau) = \bigcup_{\mathbf{p} \in entrance\ start\ line} \mathcal{R}^-(\mathbf{p}, \tau). \tag{10}$$

5.1 Bridge Assignment among Agents

We have finished describing how to build a bridge, its entrance, and the corresponding reachability regions. We will now discuss how to build a set of bridges for all the agents in the scenario, as well as how to assign a bridge to each agent. Our solution is to first build a bridge for a single agent. Next, we check whether or not there are any other agents whose initial and goal positions are within the backward and forward reachable regions of this bridge. If so, these agents can also leverage this bridge for navigation, and thus we assign this bridge to these agents. We repeatedly add more bridges until all agents have one associated bridge. The algorithm is as shown in Algorithm 4.

Algorithm 4: Generating bridges for all agents

input : Initial and goal configurations pairs $IG = \{(I_i, G_i)\}_{i=0}^{n-1}$ for all n agents
output: A set of bridges $B = \{B_j\}$

```
B ← ∅
/* compute the bridges for all of the agents                    */
while IG ≠ ∅ do
    Select two configurations (I, G) from the set IG
    Construct a 2D bridge b and its entrance using Algorithm 1
    Compute the reachability regions R+(b, τ) and R−(b, τ)
    /* compute the bridge for the agents                        */
    for (I_i, G_i) ∈ IG do
        if I_i ∈ R−(b, τ) and G_i ∈ R+(b, τ) then
            Remove (I_i, G_i) from IG
            Assign bridge b to the agent i
            Add bridge b to B
return B
```

5.2 Global Navigation using Bridges

Once the bridges and their entrances are computed, we can leverage them for efficient multi-agent global navigation. Each given agent, it will first move toward the bridge assigned to it along an optimal trajectory computed using [3]. Once it reaches the entrance, it can enter the entrance and then go through the crowded or narrow area by following the trajectory interpolated by the bridge. After leaving the bridge, the agent switches back to moving toward its individual goal following an optimal trajectory computed using [3].

6 Global Navigaion using Bridges

Once the bridges and their entrances are computed, we can leverage them for efficient multi-agent global navigation. Each given agent, it will first move toward the bridge assigned to it along an optimal trajectory computed using [3]. Once it reaches the entrance, it can enter the entrance and then go through the crowded or narrow area by following the trajectory interpolated by the bridge. After leaving the bridge, the agent switches back to moving toward its individual goal following an optimal trajectory computed using [3].

6.1 Inter-Trajectory Scheduling

However, there is still one unresolved situation: several agents may try to leverage the bridge at the same time, and this may result in collisions between the agents inside the bridge. To avoid this problem, we use a scheduling scheme among the agents so that they can share the bridge in a safe and efficient manner for collision-free navigation. This is achieved by inter-trajectory scheduling.

Algorithm 5 shows a simple scheduling algorithm among trajectories. In particular, we check whether a planned trajectory will collide with any previously scheduled trajectory. If so, we delay the trajectory for a small time δt. This process continues until all collisions among trajectories are resolved.

Algorithm 5: Inter-bridge scheduling

input : Original plans $P = \{P_i\}_{i=0}^{n}$
output: Scheduled plans $P' = \{P'_i\}_{i=0}^{n}$

```
for i = 1 to n do
    P'_i ← P_i
    for j = 0 to i − 1 do
        while P'_i collide with P'_j do
            P'_i postpone δt
return P'
```

7 Completeness and Performance Analysis

Our bridge-based planning algorithm is probabilistically complete. In particular, our trajectory generation includes two parts: the trajectory generation inside the bridge and the trajectory generation outside the bridge. The former is probabilistically complete because the bridge is computed using RRT with the start and goal positions inside the agents' reachable sets, and thus we can always find a valid bridge with probability 1 if one exists. Once the bridge is constructed, the trajectory generation inside the bridge is deterministic and thus is always possible.

The complexity of our method's online phase is $\mathcal{O}(n^2)$, where n is the number of agents. In particular, the inter-trajectory scheduling in Section 6.1 dominates the running time cost. As shown in line 5 of Algorithm 5, each agent needs to check whether its planned trajectory will collide with any plans of previously scheduled agents. If so, the agent will delay the plan to avoid collisions. For n agents, we need to execute $\frac{1}{2}n^2$ checks in total. For each check between two trajectories, we need to further check whether any two segments from these paths collide with each other. For two long trajectories P and P', seems that we need to perform a pairwise checking with an expensive $|P| \times |P'|$ complexity. However, note that when agents are moving inside the same bridge, their speeds are determined by their position in the start line and are fixed while passing through the bridge (Theorem 1). As a result, agents falling behind will never catch up the agents in the front, and thus we can determine the collision status between two plans by only checking collisions between a few segments. In fact we observe that at most $C < 10$ collision checks are performed between two trajectories in all our experiments. In this way, the computational complexity of the online phase is $C \times \frac{1}{2}n^2 = \mathcal{O}(n^2)$.

8 Experiment

	#Agent	#Bridges	# Overhead 1(ms)	# Overhead 2(ms)
Benchmark 1	96	2	204513	1012
Benchmark 2	96	2	453267	1432
Benchmark 3	100	2	553462	1431
Benchmark 4	96	2	1079381	1532

Table 1: The overhead 1 shows the time cost for constructing the given number of bridges in various benchmarks. The overhead 2 shows the overhead of the scheduling algorithm.

In this section, we describe our implementation and highlight the improved performance of our navigation algorithm on four challenging benchmarks, of which two are 3D and the others are 2D as illustrated in Figure 4, 5, 6, 7. We will

	Outside time (ms)	Inside time (ms)
Benchmark 1	23	< 1
Benchmark 2	47	< 1
Benchmark 3	132	< 1
Benchmark 4	560	< 1

Table 2: The average time cost for each agent to compute its trajectory inside and outside the bridges on various benchmarks.

	Benchmark 1	Benchmark 2	Benchmark 3	Benchmark 4
HRVO (#collision times)	-	212	-	517
HRVO+PRM (#collision times)	13	9	11	4
OURS (#collision times)	0	0	0	0

Table 3: The comparison of collision times between our method, the local method(HRVO), and the local method plus traditional global method(PRM).

	Benchmark 1	Benchmark 2	Benchmark 3	Benchmark 4
HRVO (#frames)	-	3938	-	4882
HRVO+PRM (#frames)	3905	2215	1101	752
OURS (#frames)	2024	1482	762	417

Table 4: The comparison of simulation time between our method, the local method(HRVO), and the local method plus traditional global method(PRM). The '-' marks mean agents get stuck in the obstacles so they cannot achieve their goals.

show the pre-computation overhead and running performance. In the running stage, the overhead is divided into two parts, one for single-agent trajectory calculation and the other for multi-agent trajectory collision avoidance. The algorithm was introduced in the previous section. We have implemented our algorithms in C++ on an Intel Core i7 CPU running at 3.30GHz with 16GB of RAM on Window 7. All the timings are generated on a single CPU core

In all four benchmarks, we use the secondary linear dynamic system agents. $40 - 100$ agents are introduced for each scenario, evenly distributed around the start sides of the bridges. In the 3D scenario, the radius of each agent is 3 units and its maximum velocity is 2 units/sec. The maximum acceleration is 1 $unit^2$/sec. The size of the scenario is $150 \times 80 \times 80$. In 2D the radius of the agent is 5 units. The maximum velocity of an agent is 3 units/sec. The maximum acceleration is 2 $units^2$/sec. The scenario size is 750×480 units.

We highlight our performance in several substeps of our algorithm. The results are shown in Tables 1, 2, 3, and, 4 for different benchmarks. We first analyze our online performance. In our result, the overhead for computing the global trajectory for each agents is very fast, especially when the agents are inside the bridge. In all the benchmarks, this substep of interpolating the trajectory in the bridge takes less than 1 ms. When the agents are outside the bridge, the agents can still easily find the entrance of the bridge by using the optimal trajectory equation [3]. This entrance-finding procedure introduces an overhead of less than 2 seconds for the entire simulation.

Next, we highlight the performance of our pre-computation stage. In this case, the performance depends on the implementation of the RRT algorithm. In all the experiments, these precomputation overheads are in a reasonable range. Table 1 shows the performance of our algorithm. Here the left column shows the average length of agents trajectories, i.e. how many nodes of trajectories on average. Each right column shows the number of frames. Our algorithms generate very dense motion plans. All of the agents can finish their jobs in a reasonable time. In the worst case, for 40-100 agents, our algorithm can compute global collision-free paths in 30 seconds. In order to highlight the advantages over local navigation (see Table 1), our method is compared with prior local navigation methods: RVO [5] in 3D + PRM, HRVO [6] in 2D + PRM. We describe the four scenarios and analyze the relative performance on these benchmarks.

Benchmark 1 This is a challenging scenario where 40 agents need to pass through two holes simultaneously and arrive at their respective goals, as shown in Figure 4(a). The holes are small and only allow two agents to pass through at the same time. However, the areas around the agents' start and goal positions are widely open, and thus the agents can easily find a path to enter the bridges, as shown by the small outside bridge cost in Table 1. Compared to RVO algorithms, our method can effectively compute collision-free global paths, while the RVO local navigation methods get stuck in the narrow passages and can take a long time to make all agents reach the goals. In addition, our method results in fewer collision cases as compared to RVO and RVO+PRM, as shown in Table 3.

Benchmark 2: This is a 2D benchmark with a long narrow corridor where 96 agents are trying to move from one end to the other end, as shown in Figure 5(b). The corridor only allows one agent to pass through at a time. For this scene, we compute two bridge connecting the two open spaces in the scenario. The RVO method can compute the collision-free trajectories for agents, but they tend to get stuck for a while due to the narrow corridor. The RVO+PRM method also results a large number of collisions. Our method generates a more stable and faster simulation than RVO and RVO+PRM, as shown in Figure 3.

Benchmark 3: This scenario has multiple narrow passages as shown in Figure 6(c), and we construct two bridges to help the global navigation through these narrow passages. In this scenario, the RVO method fails to compute collision-free paths because agents are getting stuck in the narrow passages. The RVO+PRM method can find a feasible solution, but it takes significantly longer time as compared to our method and also results in more collisions, as demonstrated in Table 3.

Benchmark 4: In this benchmark, we have circle obstacle with narrow space in the scene. Our method builds two bridges to connect the left and right areas in the workspace. The simulation result is shown in Figure 7(d). As compared to local methods, our approach results in fewer agent-agent collisions and can compute the final trajectories faster as shown in Table 3.

9 Conclusion and Future Work

We present a novel multi-agent global navigation algorithm using interpolation bridges. Our approach is general and overcomes some of the major limitations of prior methods in terms of navigating through crowded areas or narrow regions. We present new techniques to compute these bridges in 2D and 3D workspaces and use their properties to compute interpolating collision-free trajectories for the agents. The construction of our bridge enables collision-free multi-agent global navigation. We have demonstrated its performance on many complex 2D and 3D scenarios and can perform collision-free navigation for tens of agents in realtime.

Our approach has some limitations. The bridge computation is limited to static obstacles, or dynamic obstacles whose trajectories are known. The complexity of global navigation increases with the number of bridges in the workspace, and very complex scenarios can result in a high number of bridges. Furthermore, our current approach is limited to agents with linear dynamics. There are many avenues for future work. In addition to overcoming these limitations, we would like to design improved algorithms for bridge computation and further evaluate their performance.

References

1. Fraichard, T., Asama, H.: Inevitable collision states. a step towards safer robots? In: IEEE/RSJ International Conference on Intelligent Robots and Systems. Volume 1. (2003) 388–393 vol.1
2. Fraichard, T.: A short paper about motion safety. In: IEEE International Conference on Robotics and Automation. (2007) 1140–1145
3. Webb, D., van den Berg, J.: Kinodynamic rrt*: Asymptotically optimal motion planning for robots with linear dynamics. In: IEEE International Conference on Robotics and Automation. (2013) 5054–5061
4. Berenson, D., Abbeel, P., Goldberg, K.: A robot path planning framework that learns from experience. In: IEEE International Conference on Robotics and Automation. (2012) 3671–3678
5. van den Berg, J., Guy, S., Lin, M., Manocha, D.: Reciprocal n-body collision avoidance. In Pradalier, C., Siegwart, R., Hirzinger, G., eds.: Robotics Research. Volume 70 of Springer Tracts in Advanced Robotics. Springer Berlin Heidelberg (2011) 3–19
6. Snape, J., Guy, S., van den Berg, J., Manocha, D.: Smooth coordination and navigation for multiple differential-drive robots. In Khatib, O., Kumar, V., Sukhatme, G., eds.: Experimental Robotics. Volume 79 of Springer Tracts in Advanced Robotics. Springer Berlin Heidelberg (2014) 601–613
7. Reif, J., Sharir, M.: Motion planning in the presence of moving obstacles. In: Symposium on Foundations of Computer Science. (1985) 144–154
8. Bekris, K.E., Grady, D.K., Moll, M., Kavraki, L.E.: Safe distributed motion coordination for second-order systems with different planning cycles. The International Journal of Robotics Research **31**(2) (2012) 129–150
9. He, L., van den Berg, J.: Meso-scale planning for multi-agent navigation. In: IEEE International Conference on Robotics and Automation. (2013) 2839–2844
10. He, L., Pan, J., Wang, W., Manocha, D.: Proxemic group behaviors using reciprocal multi-agent navigation. In: IEEE International Conference on Robotics and Automation. (2016) 292–297
11. Solovey, K., Yu, J., Zamir, O., Halperin, D.: Motion planning for unlabeled discs with optimality guarantees. CoRR **abs/1504.05218** (2015)
12. Solovey, K., Halperin, D.: On the hardness of unlabeled multi-robot motion planning. CoRR **abs/1408.2260** (2014)
13. Dresner, K., Stone, P.: A multiagent approach to autonomous intersection management. Journal of Artificial Intelligence Resesearch **31**(1) (March 2008) 591–656
14. Alami, R., Fleury, S., Herrb, M., Ingrand, F., Robert, F.: Multi-robot cooperation in the martha project. IEEE Robotics Automation Magazine **5**(1) (1998) 36–47
15. Geraerts, R., Kamphuis, A., Karamouzas, I., Overmars, M.: Using the corridor map method for path planning for a large number of characters. In: Motion in Games. Springer (2008) 11–22
16. Wein, R., van den Berg, J., Halperin, D.: Planning high-quality paths and corridors amidst obstacles. The International Journal of Robotics Research **27**(11-12) (2008) 1213–1231
17. van den Berg, J., Overmars, M.: Roadmap-based motion planning in dynamic environments. IEEE Transactions on Robotics **21**(5) (2005) 885–897
18. van den Berg, J., Overmars, M.: Prioritized motion planning for multiple robots. In: IEEE/RSJ International Conference on Intelligent Robots and Systems. (2005) 430–435

19. LaValle, S., Kuffner, J.: Randomized kinodynamic planning. International Journal on Robotics Research **20**(5) (2001) 378–400
20. Karaman, S., Frazzoli, E.: Optimal kinodynamic motion planning using incremental sampling-based methods. In: IEEE Conference on Decision and Control. (2010) 7681–7687
21. Čáp, M., Novák, P., Vokrínek, J., Pěchouček, M.: Multi-agent rrt: Sampling-based cooperative pathfinding. In: International Conference on Autonomous Agents and Multi-agent Systems. (2013) 1263–1264
22. Ferguson, D., Kalra, N., Stentz, A.: Replanning with rrts. In: IEEE International Conference on Robotics and Automation. (2006) 1243–1248
23. Katsev, M., Yu, J., LaValle, S.: Efficient formation path planning on large graphs. In: IEEE International Conference on Robotics and Automation. (2013) 3606–3611
24. Sanchez, G., Latombe, J.C.: Using a prm planner to compare centralized and decoupled planning for multi-robot systems. In: IEEE International Conference on Robotics and Automation. Volume 2. (2002) 2112–2119 vol.2
25. van den Berg, J., Snoeyink, J., Lin, M., Manocha, D.: Centralized path planning for multiple robots: Optimal decoupling into sequential plans. In: Robotics: Science and systems. (2009)
26. Luna, R., Bekris, K.E.: Push and swap: Fast cooperative path-finding with completeness guarantees. In: Joint Conference on Artificial Intelligence - Volume Volume One. (2011) 294–300
27. Gottschalk, S., Lin, M.C., Manocha, D.: Obbtree: A hierarchical structure for rapid interference detection. In: SIGGRPAH. (1996) 171–180

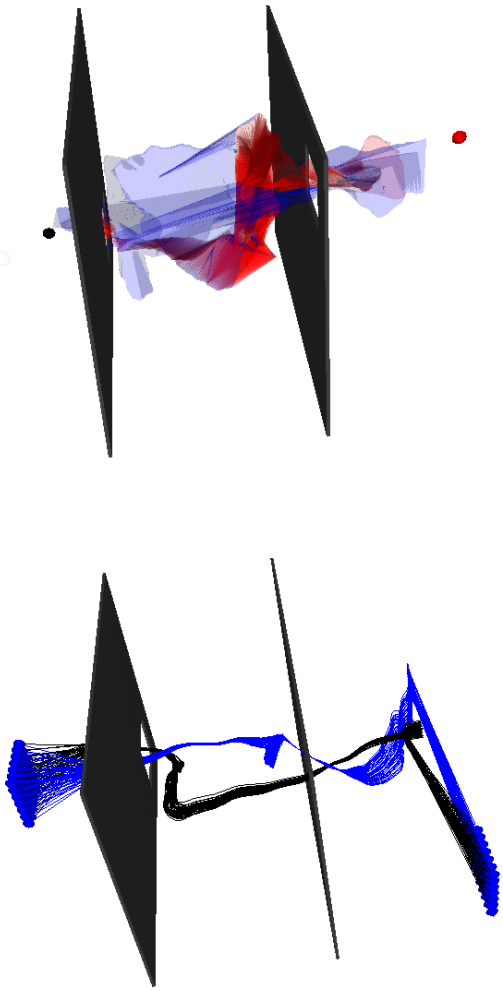

Fig. 4: This figure shows the bridges(above picture) and trajectories(below picture) of 3D scenario. In this benchmark, two groups of agents each has 49 agents exchange their positions.

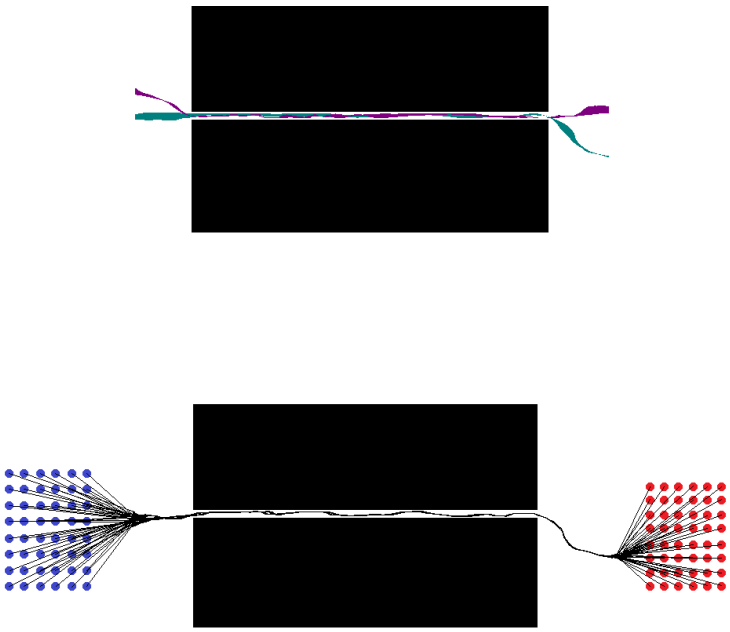

Fig. 5: This figure shows the bridges(above picture) and the trajectories(below picture) of 2D scenario 1. In this benchmark, two groups of agents each has 48 agents exchange their positions.

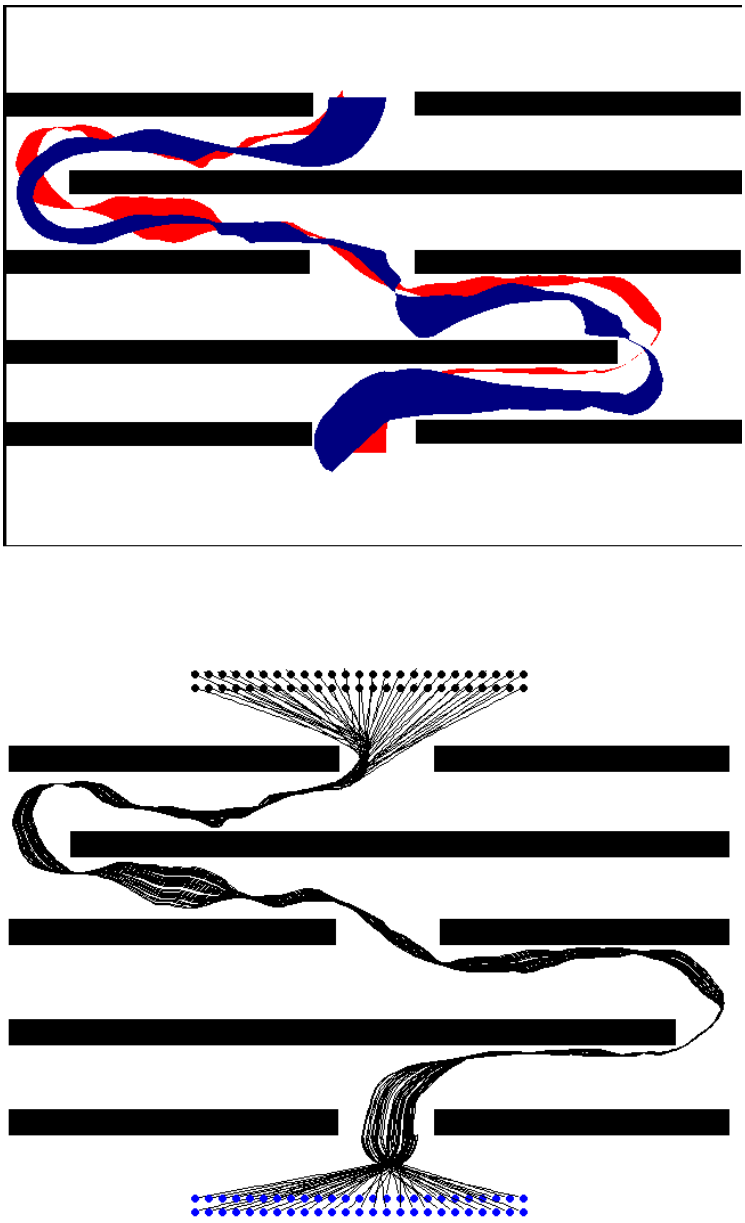

Fig. 6: This figure shows the bridges(above picture) and the trajectories(below picture) of 2D scenario 2. In this benchmark, two groups of agents each has 50 agents exchange their positions.

Fig. 7: This figure shows the bridges(above picture) and the trajectories(below picture) of 2D scenario 3. In this benchmark, two groups of agents each has 48 agents exchange their positions.